

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
23 March 2006 (23.03.2006)

PCT

(10) International Publication Number
WO 2006/030173 A1

(51) International Patent Classification:
G06K 9/64 (2006.01)

Frederick, Warwick, Michael [GB/GB]; Sheepstor,
Boyton, Woodbridge Suffolk IP12 3LH (GB).

(21) International Application Number:
PCT/GB2005/003339

(74) Agent: **LLOYD, Barry, George, Willi**; BT Group Legal
Intellectual Property Department, PP C5A, BT Centre, 81
Newgate Street, London Greater London EC1A 7AJ (GB).

(22) International Filing Date: 26 August 2005 (26.08.2005)

(81) Designated States (unless otherwise indicated, for every
kind of national protection available): AE, AG, AL, AM,
AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN,
CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI,
GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE,
KG, KM, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA,
MD, MG, MK, MN, MW, MX, MZ, NA, NG, NI, NO, NZ,
OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL,
SM, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC,
VN, YU, ZA, ZM, ZW.

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
0420727.0 17 September 2004 (17.09.2004) GB
0503944.1 25 February 2005 (25.02.2005) GB

(71) Applicant (for all designated States except US): **BRITISH
TELECOMMUNICATIONS PUBLIC LIMITED
COMPANY** [GB/GB]; 81 Newgate Street, London
Greater London EC1A 7AJ (GB).


(84) Designated States (unless otherwise indicated, for every
kind of regional protection available): ARIPO (BW, GH,
GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM,
ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM),
European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI,

(72) Inventor; and

(75) Inventor/Applicant (for US only): **STENTIFORD,**

[Continued on next page]

(54) Title: ANALYSIS OF PATTERNS



(57) Abstract: A first pattern (A) is represented by a first ordered set of elements (x_i) each having a value whilst a second pattern (B) is represented by a second ordered set of element (u_i) each having a value. The patterns are analysed by iteratively performing the steps of: selecting (e.g. at random) a plurality of elements from the first ordered set; for each selected element x_i of the first ordered set, selecting an element u_i from the second ordered set, such that the selected elements of the second ordered set have, within the second ordered set, a set of positional relationships relative to each other that is the same as, or a transformation of, the set of positional relationships that the selected plurality of elements of the first ordered set have relative to each other, comparing the value of each of the selected elements of the first ordered set with the value of the correspondingly positioned selected element of the second ordered set in accordance with a predetermined match criterion to produce a decision that the selected plurality of elements of the first ordered set does or does not match the selected plurality of elements of the second ordered set; and in the event of a match, updating at least one similarity score (S_{AB}). In the selection of the elements of the second ordered set, one chooses (e.g. at random) at least one parameter (R, β, α) and selects elements having a set of positional relationships (d_i', d_i'') that is transformed from the positional relationship set d_i of the selected elements of the first ordered set in accordance with the parameter(s). This transformation may include one or more of rotation, reflection or scaling. Other aspects include biasing the selection of the elements of the first ordered set towards a large extent, and of introducing variations in the number of elements selected, with the similarity score being updated by an amount that is a function of the number of elements. The method may also be applied to a single image.



FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, NL, PL, PT,
RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA,
GN, GQ, GW, ML, MR, NE, SN, TD, TG).

— *before the expiration of the time limit for amending the
claims and to be republished in the event of receipt of
amendments*

Published:

— *with international search report*

*For two-letter codes and other abbreviations, refer to the "Guid-
ance Notes on Codes and Abbreviations" appearing at the begin-
ning of each regular issue of the PCT Gazette.*

Analysis of Patterns

The present invention is concerned with comparing one pattern with another, or of a pattern with itself, and is of particular interest in the comparison of two-dimensional patterns such as visual images, although applicable also to one-dimensional patterns and patterns having three or more dimensions. Standard approaches to pattern recognition use templates to recognise and categorise patterns [1]. Such templates take many forms but they are normally produced by a statistical analysis of training data and matched with unseen data using a similarity measure [2]. The statistical analysis is normally carried out over a number of intuitively selected features that appear to satisfy the needs of the recognition task. For example, in speech recognition templates can be encapsulated as Hidden Markov Models derived in the frequency domain and in Optical Character Recognition the templates take the form of the character fonts themselves. In the case of face recognition a number of intuitively chosen features such as skin texture, skin colour and facial feature registration are used to define face templates [5]. In a CCTV surveillance application intruders are normally detected through a process of frame subtraction and background template modelling which detects movement and removes background effects from the processing [3]. In many cases the number of features leads to a computationally unmanageable process and Principal Components Analysis and other techniques are used to scale down the problem without significantly reducing performance [<http://www.partek.com/index.html>]. These approaches achieve great success in non-noisy environments but fail when the pattern variability and number of pattern classes increase.

Some techniques for analysis of images or other patterns where the pattern is compared with other parts of the same pattern are described in our earlier patent applications as follows.

European patent application 00301262.2 (publication No. 1126411) (applicants ref. A25904EP#);

International patent application PCT/GB01/00504 (publication No. WO 01/61648) (applicants ref. A25904WO);

International patent application PCT/GB01/03802 (publication No. WO02/21446) (applicants ref. A25055WO);

U.S patent application 977,263/09 filed 16 October 2001 (publication No. 20020081033) (applicants ref. A25904US1);

- as well as the following papers published by the inventor:

Stentiford F W M, "An estimator for visual attention through competitive novelty with application to image compression", Proc. Picture Coding Symposium 2001, Seoul, 25–27 April, pp 101 - 104, 2001.

Stentiford F W M, "An evolutionary programming approach to the simulation of visual attention",
5 Proc. Congress on Evolutionary Computation 2001, Seoul, pp 851 - 858, 27 – 30 May, 2001.

Methods and apparatus for comparing patterns are described in our earlier international patent application WO03/081523.

Aspects of the present invention are set out in the claims.

Some embodiments of the present invention will now be described, with reference to the
10 accompanying drawings, in which:

Figure 1 is a block diagram of an apparatus for performing the invention;

Figure 2 is a diagram illustrating operation of the invention;

Figure 3 is a flowchart of the steps to be performed by the apparatus of Figure 1 in
accordance with a first embodiment of the invention;

15 Figure 4 is a flowchart of the steps to be performed by the apparatus of Figure 1 in
accordance with a second embodiment of the invention; and

Figure 5 illustrates a pair of images.

Figure 1 shows an apparatus consisting of a general purpose computer programmed to
perform image analysis according to a first embodiment of the invention. It has a bus 1, to which
20 are connected a central processing unit 2, a visual display 3, a keyboard 4, a scanner 5 (or other
device, not shown) for input of images, and a memory 6.

In the memory 6 are stored an operating system 601, a program 602 for performing the
image analysis, and storage areas 603, 604 for storing two images, referred to as image A and
image B. Each image is stored as a two-dimensional array of values, each value representing the
25 brightness of a picture element within the array. It will be understood, however that the apparatus
might more generally be arranged to analyse data sets other than images, in which case the storage
areas 603, 604 would each contain a data set A, B, each being an ordered set of value, ordered in
any number of dimensions (the same number of dimensions in each set).

The image arrays are shown schematically in Figures 2a and 2b. Image A consists of an
30 array of picture elements $\underline{x}_i = (x_i, y_i)$ where x_i and y_i are the horizontal and vertical positions of the

elements within the image. A 20x20 array is shown, for the purposes of illustration. In the general case, $\underline{x}_i = (x_{i1}, x_{i2}, x_{i3}, \dots, x_{in})$ where x_{ij} are the n coordinates of \underline{x}_i in n dimensions. Each element has a respective value $a = a(\underline{x}_i)$. This may be a scalar value a or vector (multidimensional) value \mathbf{a} . In the case of an image these would be the brightness a , or alternatively a set of colour components such as (in r,g,b representation) $\mathbf{a} = (a_r, a_g, a_b)$. Similarly, the image B consists of an array of picture elements $\underline{u}_i = (u_i, v_i)$ having brightness values $b(\underline{u}_i)$. In the general case, $\underline{u}_i = (u_{i1}, u_{i2}, u_{i3}, \dots, u_{in})$.

Figure 3 is a flowchart explaining the application of the method as applied to two-dimensional images A, B. In Step 200, a score S_{AB} is set to zero, and (201) a counter t set to zero.

The first task of the method is to make a random selection of elements in image A. This may be visualised as stabbing the image with an irregular fork having m tines. The actual number m of elements may be selected (Step 202) at random within the range $m_{\min} \leq m \leq m_{\max}$, or m may be fixed, in which case Step 202 would be omitted. The fork size is limited by a diameter D selected as a random number $D_{\min} \leq D \leq D_{\max}$ (Step 203). Other approaches to varying the distribution of fork size will be discussed later. Note that, throughout this specification, references to a random selection also envisage the possibility of selection by means of a pseudo-random process. Also, a reference to random selection envisages the possibility not only of a random selection for which the selection of any value (within a range specified) is equally probable, but also of a selection in accordance with a non-uniform probability distribution.

At Step 204 a set N_x of m elements at random positions \mathbf{x}_i are selected, subject to the constraint that all elements lie within the image and all elements lie within a circular area of diameter D . One possible method of doing this is as follows:

2041 Choose a random element \mathbf{x}_0 in A.

2042 Choose a set of random displacements $\mathbf{d}_i = (d_{xi}, d_{yi})$, where $0 \leq d_{xi} \leq D/2$ and $0 \leq d_{yi} \leq D/2$.

2043 If the Euclidean distance $|\mathbf{d}_i| \leq D/2$, then continue; otherwise repeat the previous step until this condition is satisfied. Where $|\mathbf{d}_i| = \sqrt{d_{xi}^2 + d_{yi}^2}$.

2044 Compute positions $\mathbf{x}_i = \mathbf{x}_0 + \mathbf{d}_i$

2045 If any \mathbf{x}_i lies outside the image area, repeat steps 2042 to 2045 in respect of the offending \mathbf{d}_i .

Note that this process permits the occurrence of one or more pairs $\mathbf{x}_i = \mathbf{x}_j$ $i \neq j$ i.e. the \mathbf{x}_i are not constrained to be distinct. This is not objectionable in practice but if desired may be

eliminated by testing for it and reselecting one member of the matching pair in the same way as for elements falling outside the image area..

The next stage (Step 205) is to make a selection N_u of m elements u_i within image B, at a random overall position, but having the same positions relative to one another as the selected elements of image A. That is, $x_i - x_j = u_i - u_j$ for all i, j . This may be visualised as stabbing image B with the same fork as used on image A.

Assuming that a set of displacements d_i have been calculated as above, this can be done by:

- 2051 Choose a random element u_0 in B.
- 10 2052 Compute positions $u_i = u_0 + d_i$
- 2053 If any u_i lies outside the image area, repeat steps 2051 to 2043 in respect of the offending d_i .

At step 206 it is determined whether each of the selected elements in image A matches the element having the same relative position in image B. Where the value associated with each element is a scalar value, in this case representing brightness, the test is that a match is deemed to occur if

$$|a(x_i) - b(u_i)| < \delta \quad \text{for all } i = 1 \dots m$$

where δ is some small threshold value

In the vector case (e.g. r, g, b) a match occurs if

$$20 \quad \text{Dist}[a(x_i) - b(u_i)] < \delta \quad \text{for all } i = 1 \dots m$$

Where Dist is some distance operator (e.g. Euclidean or city-block), or the vector components may be thresholded separately, e.g. for colour components

$$\begin{aligned} &|b_r(x_i) - b_r(u_i)| < \delta_r \quad \text{and} \\ &|b_g(x_i) - b_g(u_i)| < \delta_g \quad \text{and} \\ 25 \quad &|b_b(x_i) - b_b(u_i)| < \delta_b \quad \text{for all } i = 1 \dots m. \end{aligned}$$

If a match occurs, then at Step 207 the Score S_{AB} is incremented. It could simply be incremented by 1 (or other fixed value). Where m is variable, the score could if desired be incremented by m or some function of m (to be discussed below).

Once the score has been updated, or if no match occurred, then the iteration counter t is tested at 209 to see if it has reached a maximum iteration count T , and if so, the process terminates. Otherwise the process returns to Step 202 for a further iteration.

A discussion of fork diameter is in order. If small forks are used, then the score or similarity measure will depend on local similarities and will assume high values if one image possesses many regions that are similar to those in the other. Large diameters give more weight to larger-scale similarities between images. If D is fixed, then one has a certain statistical distribution of actual fork diameters ranging from very small up to the maximum permitted by the particular value of D (by actual diameter we mean the diameter of the smallest circle that can enclose all the elements x_i on a particular iteration). The random selection of D in the range $D_{\min} \leq D \leq D_{\max}$ allows actual diameters up to D_{\max} but skews their statistical distribution by favouring actual diameters less than D_{\min} at the expense of diameters between D_{\min} and D_{\max} .

Alternative methods of skewing this distribution can be envisaged (including the possibility of favouring larger forks); for example one could proceed with a fixed value of D , then, for each selection of a set of d_i ($i = 1, \dots, m$) one could calculate the actual diameter (or an estimate of it such as $\sqrt{\{[\text{Max}(d_{xi}) - \text{Min}(d_{xi})]^2 + [\text{Max}(d_{yi}) - \text{Min}(d_{yi})]^2\}}$) and discard the set or not with a probability corresponding to a desired function of the actual diameter.

Turning back to the score-incrementing step, in the case of variable m , the information value of a match rises according to its rarity. Thus a match obtained using a fork with many tines is more significant than one obtained using few tines. Hence the suggestion above that the match be incremented by m . More generally, one might increment the score by some monotonically increasing function $f(m)$ of m rather than m itself. Functions that rise more rapidly than m itself may for example be used (such as m^2 or 2^m).

We turn now to a second embodiment of the invention, which introduced a number of additional transformations of the fork applied to image B. The following description omits the random variation of m and D discussed above, though these could be included (or not) if desired.

The idea here is to allow matches to be obtained, and hence contributions to the total score to be generated, in situations where a feature appears in one image whilst a similar feature appears in the other image but is of a different size, is reversed, or is rotated with respect to the one in the first image. This is achieved by applying to image B a fork that is not the same as that applied to image A but is a scaled and/or reversed and/or rotated version of it.

Figure 4 shows a flowchart. Steps that are identical to those in Figure 3 are given the same reference numerals. Steps 2041 to 2045 are not shown explicitly.

Following Step 204, a random reflection is applied to the fork:

Choose a random integer R as 0 or 1;

If $R = 1$, set $d_{xi} = -d_{xi}$ for all $i = 1 \dots m$. d_{yi} are unchanged.

This is a reflection about a single vertical axis. Reflections about other axes are encompassed by this reflection plus the rotation step that comes next.

Next (Step 211) choose a random rotation angle β in the range β_{\max} to β_{\min} . Then (212) the
5 fork represented by d_i is rotated by this angle by performing, for each i :

$$d_{xi}' = d_{xi} \cdot \cos\beta - d_{yi} \cdot \sin\beta$$

$$d_{yi}' = d_{xi} \cdot \sin\beta + d_{yi} \cdot \cos\beta.$$

Thirdly, in Step 213, a random scale factor α is chosen in the range α_{\min} to α_{\max} . The fork
is then scaled at 214 by this factor:

$$10 \quad d_{xi}'' = \alpha \cdot d_{xi}'$$

$$d_{yi}'' = \alpha \cdot d_{yi}'$$

The scaling and rotation are relative to (0,0), though scaling or rotation about any other centre would do just as well.

If desired, it would be possible to define α_x , α_y separately in x and y directions, for
15 example to spot longer or taller versions of the same shapes.

Step 215 then follows to determine the elements u_i . This is identical to Step 205 except that it uses the transformed d_i'' rather than d_i . Note that there is a possibility that the selection of the transformations, or even of d_i , may mean that it is difficult or even impossible to find, in Step 215 (which comprises the substeps 2051 to 2053 described earlier), a position u_0 such that all
20 elements u_i lie within image B. Thus, although not shown in the flowchart, it may be necessary in such an event to return to step 210 for a fresh selection of the transformation(s) or even return to step 204 for an entirely new fork.

Steps 206 to 209 are as described earlier.

By way of illustration, Figure 5 shows the two dimensional case in which a fork with $m =$
25 4 pixels x in image A is matched with 4 pixels u in image B after a 90° rotation and a 25% scale reduction. Each of the pixels might possess three colour intensities, so $a = (a_r, a_g, a_b)$ and the pixels u match the x if the colour intensities of all m corresponding pixels have values within δ_j of each other.

Note that the reflection, rotation and scaling are linear operations and can be performed in
30 any order. The use of these transforms in addition to translation means that structural similarities between the two images will not be obscured by reflection, rotation or differences in scale.

It is not essential to employ all three transformations. One may be used alone, or two in combination. It is probably rare that would want to allow reflection but not allow rotation, though it would be possible. If reflection about any axis were required, but not rotation, then either the reflection step could be modified to reflect about a randomly chosen axis, or (which amounts to the same thing) the rotation step could be retained but skipped when $R=0$. One might choose to allow scaling but not rotation in cases where similarities invariant to scale but not rotation are being sought - e.g. it might be required that diamond shapes are to be distinguished from squares.

In a further modification, values of parameters that result in a match may be recorded. A simple implementation of this would be, for example to record each rotation angle β that results in a match. These values could be averaged. For example an average β of 22° could be useful as indicating that image B is rotated 20° relative to image A. A spread in angles for each match would indicate a measure of circular symmetry present in both images. E.g. matches on a red ball in images A and B could occur for any angular rotation.

Similar processing might be applied to the scale factors α .

Another example is to compute, in the case of a match the position of the centroid g_x of x_i and the centroid g_u of u_i . Averaging g_x would give the position in image A of a region that is similar to a region at a position in image B given by the average of g_u . In practice, since the centroid g is the average of x_i (or u_i) then (if m constant) there is no need to calculate g and then average them; one could just average all the x_i (or u_i). Again, this approach gives useful results only if the matches mainly result from one region in each image. A more sophisticated approach is to determine, for each match, whether the positions of the centroids g_x and g_u are similar to the centroid pair recorded for a previous match. E.g.

$\text{Dist}(g_x - g_x') < e$ and $\text{Dist}(g_u - g_u') < e$ (where Dist is the Euclidean distance or other distance operator and g_x', g_u' are the pair of centroids corresponding to an earlier match (or conceivably a running average of such earlier matched centroids) and e is some small threshold). This would imply that the match involved the same region as did the earlier match. One could then proceed by including the new result into the average only if this condition were satisfied. More than one average could be computed, for more than one different region, but to do this it would probably be more efficient to store the data for each match and process it at the end.

The set of region correspondences may be used to track objects between video frames or obtain non-linear registration between images (e.g. the registration of mammograms taken at different times). If image A is simultaneously compared with multiple images B, C, D, ..., Z, the

similarity measure may be used to detect the presence of specific objects or compare the state of scenes with a set of normal and abnormal states (e.g. for security purposes). In a similar fashion the system may be used to identify facial expressions or facial identities.

The same kind of averaging, conditional upon similarity of centroid positions could also be applied to the rotation angle β . Such a process would indicate a strong similarity between a square and a diamond but only after a rotational transform of 45° , 135° , 225° , or 315° .

Typical values for the various parameters given above, for 600x400 images, are as follows

m_{\min}	1
m_{\max}	20
D_{\min}	1
D_{\max}	A smaller value may be chosen for extracting region similarities, larger for global similarities. May range from 10 up to the smallest of the widths and heights of the two images
δ	30 (assuming that a lies in the range 0 to 255)
$\delta r, \delta g, \delta b$	30 (assuming that r, g, b lie in the range 0 to 255)
T	20000
β_{\min}	-90
β_{\max}	+90
α_{\min}	0.5
α_{\max}	2
e	10

In the case of images A, B both containing large plain areas of similar brightness (or colour), a large number of matches and hence a high score might be obtained even though there may be little similarity between the parts of the image that contain detail. In this case, the process need not necessarily be carried out for the whole image. For example, if regions of images A and B have been identified as being of special interest – perhaps using the method described in one of our earlier patent applications referred to above – then the picture elements \underline{x} , \underline{u} dealt with may be just those lying in the vicinity of the identified region. This could be implemented by constraining the selection of \underline{x} and \underline{u} such that at least one element x_i is an element of a special interest region of image A and at least one element u_i is an element of a special interest region of image B. This means that each match establishes a similar relationship in each image between the foreground and background, or foreground and foreground (if all times lie on high interest pixels). Small diameter forks will tend to obtain similarities between localised regions in A and B, and larger forks will determine global relationships between foreground and background.

The methods described in two-dimensions can equally well be applied to one dimensional data such as temporal sequences of values (e.g. samples of an audio or other signal). However, rotation has no meaning in one dimension. They may likewise be applied in three (or more)

dimensions, where, however rotation requires more angles (up to $n-1$) to be specified if rotation is to be permitted about more than one axis. Rotation has meaning only if the disturbed dimensions are in the same category. Thus a three dimensional set of values representing a sequence of visual images (x, y, t) (or a four -dimensional set representing a temporal sequence of three-dimensional visual images (x, y, z, t)) could be rotated only about axes parallel to the time axis.

Discussion

This method carries out an analysis of individual pairs of images and extracts structure that is common to both images. The amount of structure that is common to both images gives a measure of their similarity. Unlike other methods the approach makes no use of pre-defined feature measurements and it does not require representative training sets of images and associated processing to optimise recognition parameters.

Scale and Orientation Differences

Scale and orientation differences in the visual domain can arise because of perspective, viewing angle and other factors and it is common for prior knowledge to be incorporated in the classifier to compensate for such differences. Such prior knowledge is not necessary in this method because the matching process takes account of differences arising from object orientation and scale. This means, for example, that the method can be applied to recognition problems in which the object to be recognised is embedded in another image. This method therefore also has application to the problem of detecting copyright infringement where portions of material have been cropped from larger works (of art, for example), and the task of reducing the size of databases where it is known that duplication is prevalent.

Local Distortions

Unless provision is made for specific pattern content, the standard template approaches to pattern recognition fail when the patterns under comparison differ because of local distortions, or small movements, as would be the visual case with rustling trees, moving clouds, changes in facial expression, scene-of-crime fingerprints on irregular surfaces, or noise, for example. Such provision requires prior knowledge of the application and will still cause the system to fail if the unseen pattern distortions do not conform to the system design requirements. This method is able to minimise the effects of local distortions without prior knowledge of the type of distortion.

The method is also able to cope with partial patterns that contain gaps and omissions. It therefore has direct application to the problem of identifying scene-of-crime fingerprints where only a part of the total print is available for matching. In the case of facial recognition the method

lends itself to searches based upon restricted portions of the unknown face. This means, for example, that searches can be carried out purely on the basis of the eye and nose region in cases where beards and moustaches might lead to ambiguity. However, the performance will degrade steadily as the distortions and gaps increase, but will not suddenly fail catastrophically.

5 *Image Classification*

The method may be used to calculate similarity scores between an image and a set of reference patterns. Similarity scores can be used to cluster groups of such patterns possessing high similarity scores relative to each other. Vantage patterns to which other patterns in the cluster are similar, taken from each cluster may themselves be clustered to form super-clusters and the
10 process continued to structure very large pattern databases. Query-By-Example retrieval is carried out by measuring similarity scores to each of the top level vantage patterns and then to each of the vantage patterns in sub-clusters corresponding to the highest scoring vantage pattern in the previous cluster. It is likely that some vantage patterns will represent clusters with common elements in those cases where patterns possess high similarity scores with more than one vantage
15 pattern. Patterns that may be handled in this way include faces and facial expressions. It may be applied to the categorisation of images of manufactured or processed materials such as sand and aggregates). In a military context a flying object might be detected and values of similarity scores would reveal whether the object was likely to be a bird or a plane.

The method may be applied to patterns of any dimension, such as one-dimensional audio
20 signals, three dimensional video data (x,y,time), or n-dimensional time dependent vectors derived from any source such as sensor arrays. In the case of speech recognition the method is able to handle variations in the speed of the speech without the use of special heuristics. A conventional approach uses Dynamic Time Warping to overcome this problem, but invokes greater computational effort and the danger of the warping process leading to increased misclassifications
25 especially in a large multi-class problem. Furthermore portions of utterances would be sufficient for word identification using this method if they were unique in the domain of discourse (e.g. 'yeah' instead of 'yes', missing the final sibilant).

Disparity Processing

The method may be applied to the problem of disparity detection as is the case when
30 detecting motion or parallax. Standard approaches to disparity detection rely heavily upon accurate registration between two images so that the subtraction (which may be carried out piecewise for a small areas of the total image) takes place between pixels that correspond to the

same points on the original object pictured in the two images. The resulting difference-image highlights those areas that correspond to differences in the original images. This becomes extremely difficult if noise is present as uncertainty is introduced into the estimate of the correct registration position and many spurious differences can be generated as a result. Even if noise is
5 absent local distortions or slight subject movements will cause mis-registration and areas of difference will be highlighted which are of little interest unless the distortion or movement itself is being measured. Linear or non-linear digital image registration techniques prior to subtraction partially compensates but does not eliminate this problem in a large proportion of cases [4]. This method obtains registration by locating correspondences between similar regions in a pair of
10 images using features that are present in both regions; those regions that do not match are ignored. Such correspondences between regions in successive frames of a video can be applied to the problem of recognising and tracking moving objects.

This method also has the advantage of being able to detect multiple disparities in which image A is compared with images B, C, D, etc. This would be useful in the case of CCTV intruder
15 detection in which image frames B, C, D, etc would be typical examples of different atmospheric conditions and other normal background states, and an alarm would be raised only if disparities were detected in all the normal image frames. It would also be applicable to the problem of measuring the likeness of facial images.

Single Image

20 Note that the method can be applied in the case where image A and image B are identical. Indeed, if desired, one could store a single image A in storage area 603 and storage area 604 becomes redundant (and references above to image B are replaced by image A).

In the case when image A is identical to image B, and the process is in effect considering a single image, a high similarity score (S_{AA}) may arise from application of the reflection transform if
25 symmetry is present in the image across the axis of reflection. In a similar fashion rotational symmetry and the presence of perspective symmetries may be detected. The distributions of rotation and reflection axis angles at which matches are found indicate the orientations of the symmetries present in the image. Forks should include some high attention scoring pixels otherwise large background tracts of self-matching sky, for example, would appear to exhibit
30 trivial symmetries. This could be achieved by restricting the process to one or more areas identified as being of special interest, as discussed earlier.

Hardware Implementation

The algorithm is eminently suitable for parallel implementation as the processing for each pixel fork is independent of the processing of other forks. This means that processing of forks may be distributed across many independent sources of computation thereby obtaining processing speeds that are only limited by the rate of data capture.

Some image analysis techniques carry out comparison calculations between images using patches that are forks in which all the pixels are employed. Patches match when a measure of correlation exceeds a certain threshold. These approaches are unable to make best use of detailed structure that is smaller than the size of the patch except in the case in which the correlation measure is designed to identify a specific texture. In addition such patches are of a defined size; if they are too large they fail to match anywhere; if they are too small they match in the wrong places. The sparse pixel forks N_x used in this method do not suffer from these constraints.

References

- [1] Vailaya A. et al., Image Classification for Content-Based Indexing, IEEE Trans on Image Processing, Vol 10, No 1, pp 117 - 130, Jan 2001.
- [2] Santini S. & Jain R., Similarity Matching, in Proc 2nd Asian Conf on Computer Vision, pages II 544-548, IEEE, 1995.
- [3] IEEE Trans PAMI - Special Section on Video Surveillance, vol 22 No 8, Aug 2000.
- [4] Brown L. G., A survey of image registration techniques, ACM Computing Surveys, Vol.24, No. 4 (Dec. 1992), pp. 325-376.
- [5] Zhao W., Chellappa R., Rosenfeld A., and Phillips P.J., "*Face recognition: A literature survey.*" CVL Technical Report, University of Maryland, October 2000.
<<ftp://ftp.cfar.umd.edu/TRs/CVL-Reports-2000/TR4167-zhao.ps.gz>>.
- [6] Sebastian T.B., Klein P.N., and Kimia B.B., "Recognition of shapes by editing shock graphs," Proc ICCV 2001, pp 755-762. Source of data at <http://www.lems.brown.edu/vision/researchAreas/SIID/>

Claims

1. A method of comparing a first pattern (A) represented by a first ordered set of elements (x_i) each having a value with a second pattern (B) represented by a second ordered set of element (u_i) each having a value, comprising iteratively performing the steps of:

- 5 (i) selecting a plurality of elements from the first ordered set ;
- (ii) for each selected element of the first ordered set, selecting an element from the second ordered set, such that the selected elements of the second ordered set have, within the second ordered set, a set of positional relationships relative to each other that is the same as, or a transformation of, the set of positional relationships that the selected plurality of elements of the
- 10 first ordered set have relative to each other,

- (iii) comparing the value of each of the selected elements of the first ordered set with the value of the correspondingly positioned selected element of the second ordered set in accordance with a predetermined match criterion to produce a decision that the selected plurality of elements of the first ordered set does or does not match the selected plurality of elements of the second
- 15 ordered set;

(iv) in the event of a match, updating at least one similarity score (S_{AB});

- wherein the step of selecting the elements of the second ordered set comprises choosing at least one parameter (R, β, α) and selecting elements having a set of positional relationships (d_i', d_i'') that is transformed from the positional relationship set (d_i) of the selected elements of the first ordered
- 20 set in accordance with the parameter(s).

2. A method of analysing a pattern (A) represented by an ordered set of elements (x_i) each having a value, comprising iteratively performing the steps of:

- (i) selecting a first plurality of elements from the ordered set ;
- 25 (ii) for each selected element of the ordered set, selecting a further element from the ordered set, such that the selected further elements of the ordered set have, within the ordered set, a set of positional relationships relative to each other that is the same as, or a transformation of, the set of positional relationships that the selected first plurality of elements of the ordered set have relative to each other,

(iii) comparing the value of each of the first selected elements of the ordered set with the value of the correspondingly positioned selected further element of the ordered set in accordance with a predetermined match criterion to produce a decision that the selected first plurality of elements of the ordered set does or does not match the selected further plurality of elements of the ordered set;

(iv) in the event of a match, updating at least one similarity score (S_{AA});

wherein the step of selecting the further elements of the ordered set comprises choosing at least one parameter (R, β, α) and selecting elements having a set of positional relationships (d_i', d_i'') that is transformed from the positional relationship set (d_i) of the selected first elements of the ordered set in accordance with the parameter(s).

3. A method according to claim 1 or 2 in which the selection of a plurality of elements from the first ordered set from the ordered set is random or pseudo-random.

4. A method according to claim 1, 2 or 3 in which the selection of the parameter(s) is random or pseudo-random.

5. A method according to any one of the preceding claims in which the or a said parameter (R) determines whether or not a reflection should be applied to the positional relationships.

6. A method according to any one of the preceding claims in which the or a said parameter (β) determines whether and the degree to which a rotation is to be applied to the positional relationships.

7. A method according to any one of the preceding claims in which the or a said parameter (α) determines whether and the degree to which a scaling is to be applied to the positional relationships.

8. A method according to any one of the preceding claims including analysing those values of the parameters that result in a match.

9. A method according to any one of the preceding claims wherein the selections of elements in the first ordered set is such that their extent, within the ordered set, is biased towards a small extent rather than a large extent.

5

10. A method of comparing a first pattern represented by a first ordered set of elements each having a value with a second pattern represented by a second ordered set of element each having a value, comprising iteratively performing the steps of:

(i) selecting a plurality of elements from the first ordered set ;

10 (ii) for each selected element of the first ordered set, selecting an element from the second ordered set, such that the selected elements of the second ordered set have, within the second ordered set, a set of positional relationships relative to each other that is the same as, or a transformation of, the set of positional relationships that the selected plurality of elements of the first ordered set have relative to each other,

15 (iii) comparing the value of each of the selected elements of the first ordered set with the value of the correspondingly positioned selected element of the second ordered set in accordance with a predetermined match criterion to produce a decision that the selected plurality of elements of the first ordered set does or does not match the selected plurality of elements of the second ordered set;

20 (iv) in the event of a match, updating at least one similarity score;

wherein the selections of elements in the first ordered set is such that their extent, within the ordered set, is biased towards a small extent rather than a large extent.

11 A method according to claim 9 or 10 comprising defining a bounded region of the first
25 ordered set, said bounded region having a randomly selected size (D), and constraining selection of the elements such that the selected elements lie entirely within the defined bounded region.

12. A method according to any one of the preceding claims in which, on each iteration, the number of elements (m) to be selected from the first ordered set is selected at random, and in

which the step of updating the score increments the score by an amount which is a monotonically increasing function of the number (m) of elements.

13. A method of comparing a first pattern represented by a first ordered set of elements each
5 having a value with a second pattern represented by a second ordered set of element each having a value, comprising iteratively performing the steps of:

(i) selecting a plurality of elements from the first ordered set ;

(ii) for each selected element of the first ordered set, selecting an element from the second
ordered set, such that the selected elements of the second ordered set have, within the second
10 ordered set, a set of positional relationships relative to each other that is the same as, or a transformation of, the set of positional relationships that the selected plurality of elements of the first ordered set have relative to each other,

(iii) comparing the value of each of the selected elements of the first ordered set with the
value of the correspondingly positioned selected element of the second ordered set in accordance
15 with a predetermined match criterion to produce a decision that the selected plurality of elements of the first ordered set does or does not match the selected plurality of elements of the second ordered set;

(iv) in the event of a match, updating at least one similarity score;

in which, on each iteration, the number of elements (m) to be selected from the first ordered set is
20 selected at random, and in which the step of updating the score increments the score by an amount which is a monotonically increasing function ($f(m)$) of the number (m) of elements.

14. A method according to claim 12 or 13 in which the function of the number (m) of elements
is equal to the number of elements.

25

15. A method according to claim 12 or 13 in which the function increases more rapidly than
the number (m) of elements than does the number (m) of elements.

16. A method according to any one of the preceding claims, in which each value (a) comprises
30 a plurality of components.

17. A method according to any one of the preceding claims, comprising firstly processing the first and second ordered set to identify regions thereof containing significant detail, and in which the selection of elements is constrained such that at least one of the selected elements of the first
5 ordered set shall lie in the or an identified region of the first ordered set and/or that at least one of the selected elements of the first ordered set shall lie in the or an identified region of the first ordered set.
18. A method according to any one of the preceding claims including analysing those values
10 of the selected element positions that result in a match.

1/5




Fig. 1

2/5




Fig. 2a




Fig. 2b

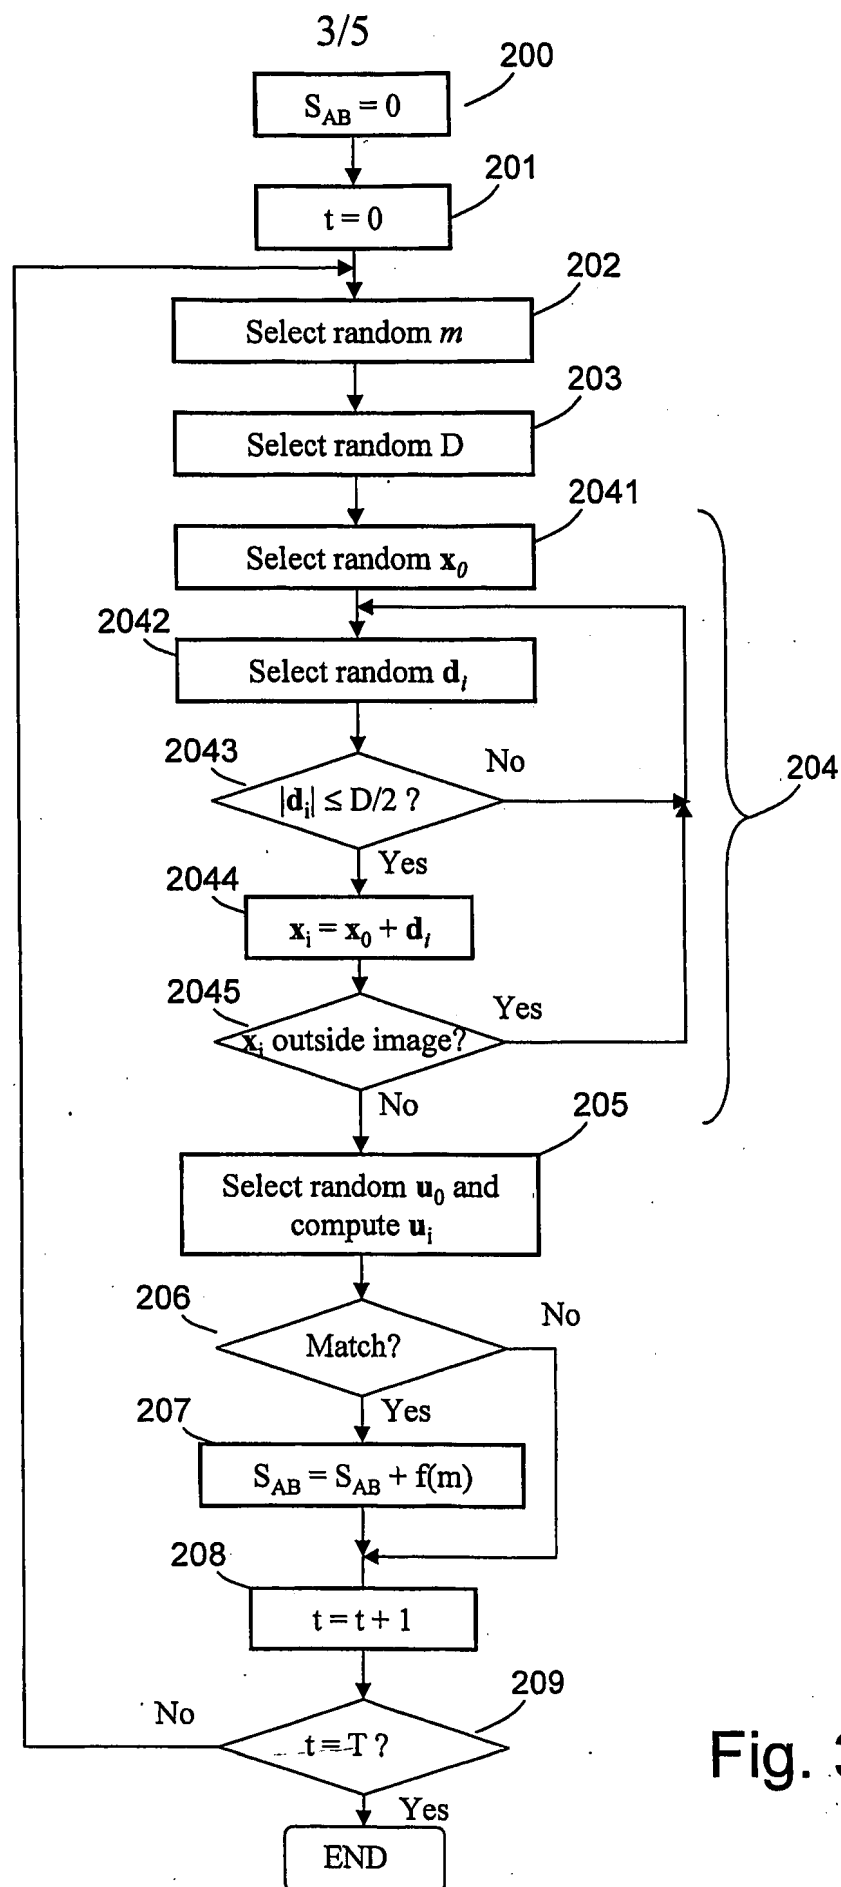


Fig. 3




Fig. 4.

BEST AVAILABLE COPY

5/5




Fig. 5

INTERNATIONAL SEARCH REPORT

Int'l application No

PCT/GB2005/003339

A. CLASSIFICATION OF SUBJECT MATTER
G06K9/64**BEST AVAILABLE COPY**

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
G06K

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, INSPEC

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	WO 03/081523 A (BRITISH TELECOMMUNICATIONS PUBLIC LIMITED COMPANY; STENTIFORD, FREDERI) 2 October 2003 (2003-10-02) cited in the application the whole document	
A	STENTIFORD F: "An attention similarity measure for fingerprint retrieval" PROC. 4TH EUROPEAN WORKSHOP ON IMAGE ANALYSIS FOR MULTIMEDIA INTERACTIVE SERVICES, 9 April 2003 (2003-04-09), - 11 April 2003 (2003-04-11) pages 27-30, XP002348928 page 2 - page 3 ----- -/-	

☒ Further documents are listed in the continuation of Box C.☒ See patent family annex.

* Special categories of cited documents:

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

- "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- "&" document member of the same patent family

Date of the actual completion of the international search

13 October 2005

Date of mailing of the international search report

09.02.2006

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Müller, M

INTERNATIONAL SEARCH REPORT

International application No

PCT/GB2005/003339

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	WO 99/60517 A (DATACUBE, INC) 25 November 1999 (1999-11-25) abstract page 2, paragraph 1 - paragraph 2 page 4, paragraph 2 - page 6, paragraph 3 -----	
A	CHANG S ET AL: "Fast algorithm for point pattern matching: invariant to translations, rotations and scale changes" PATTERN RECOGNITION, ELSEVIER, KIDLINGTON, GB, vol. 30, no. 2, February 1997 (1997-02), pages 311-320, XP004056826 ISSN: 0031-3203 abstract page 311, left-hand column - page 312, right-hand column, paragraph 4 -----	

BEST AVAILABLE COPY

INTERNATIONAL SEARCH REPORT

ational application No.
PCT/GB2005/003339

Box No. II Observations where certain claims were found unsearchable (Continuation of Item 2 of first sheet)

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:
2. ☐ Claims Nos.:
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:
3. ☐ Claims Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box No. III Observations where unity of invention is lacking (Continuation of Item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

see additional sheet

1. ☐ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. ☐ As all searchable claims could be searched without effort justifying an additional fees, this Authority did not invite payment of additional fees.
3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:
4. ☒ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

1-9 (+11, 12 and 14-18 insofar as dependent on claims 1-9)

Remark on Protest

- ☐ The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.
- ☐ The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.
- ☐ No protest accompanied the payment of additional search fees.

FURTHER INFORMATION CONTINUED FROM PCT/ISA/ 210

This International Searching Authority found multiple (groups of) inventions in this international application, as follows:

1. claims: 1-9 (+ 11, 12 and 14-18 insofar as dependent on claims 1-9)

Pattern matching based on randomized point selection which is robust to reflection, rotation or scaling

2. claims: 10-12

Pattern matching based on randomized point selection the extent of which is biased to be rather smaller than larger.

3. claims: 13-18

Pattern matching based on randomized point selection where the matching score increases monotonically with the number of points to be matched.

INTERNATIONAL SEARCH REPORT

International application No

PCT/GB2005/003339

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
WO 03081523	A	02-10-2003	AU 2003215755 A1	08-10-2003
			CA 2479223 A1	02-10-2003
			CN 1643540 A	20-07-2005
			JP 2005521165 T	14-07-2005
			US 2005169535 A1	04-08-2005

WO 9960517	A	25-11-1999	CA 2332118 A1	25-11-1999
			CN 1301369 A	27-06-2001
			EP 1080444 A1	07-03-2001
			JP 2002516440 T	04-06-2002

BEST AVAILABLE COPY